

THE BELL SYSTEM TECHNICAL JOURNAL

VOLUME XLI

SEPTEMBER 1962

NUMBER 5

Copyright 1962, American Telephone and Telegraph Company

On Rearrangeable Three-Stage Connecting Networks

By V. E. BENEŠ

(Manuscript received December 26, 1961)

A class of three-stage connecting networks proven rearrangeable by D. Slepian is considered. Bounds on the number of calls that must be moved are obtained by some simple new methods.

1. INTRODUCTION

Most communications systems contain a *connecting network* as a basic functional unit. A connecting network is an arrangement of switches and transmission links through which certain terminals can be connected together in many combinations.

The calls in progress in a connecting network do not usually arise in a predetermined time sequence. Requests for connection (new calls) and terminations of connection (hangups) occur more or less at random. For this reason the performance of a connecting network when subjected to random traffic is used as a figure of merit. This performance is measured, for example, by the fraction of requested connections that cannot be completed, or the *probability of blocking*.

The performance of a connecting network for a given level of offered traffic is determined largely by its *configuration or structure*. This structure may be described by stating what terminals have a switch placed between them, and can be connected together by closing the switch. The structure of a connecting network determines what combinations of terminals can be connected together simultaneously. If this structure is

too simple, only a few calls can be in progress at the same time; if the structure is extensive and complex, it may indeed provide for many large groups of simultaneous calls in progress, but the network itself may be expensive to build and difficult to control.

The structure of a connecting network also gives rise to various purely combinatory properties that are useful in assessing performance. For example, C. Clos¹ has exhibited a whole class of connecting networks that are *nonblocking*: no matter in what state the network may be, it is always possible to connect together an idle pair of terminals without disturbing calls already in progress. We call such a network *nonblocking in the strict sense*, because it has no blocking states whatever.

If a connecting network does have blocking states, it is nevertheless possible that by suitably choosing routes for new calls one can confine the trajectory of the operating system to nonblocking states. That is, there may exist a rule whose use in putting up new calls results in avoiding all the blocking states, so that the system is effectively nonblocking. The rule only affects new calls that could be put into the network in more than one way; no call already in progress is to be disturbed. Connecting networks for which such a rule exists we call *nonblocking in the wide sense*.

In practice, the procedure of routing the calls through the network is called "packing" (the calls), and the method used to choose the routes is called a "packing rule." The use of the word "packing" in this context was undoubtedly suggested by a natural analogy with packing objects in a container. Virtually nothing rigorous is known about the effect of packing rules on network performance.

Finally, a connecting network may or may not have the property of being *rearrangeable*: given any set of calls in progress and any pair of idle terminals, the existing calls can be reassigned new routes (if necessary) so as to make it possible to connect the idle pair.

These three combinatory properties of connecting networks have been given general topological characterizations in a previous paper.² In this paper we consider the last property described, that of rearrangeability, and we study the extent to which it applies to a specific class of connecting networks.

Fig. 1 shows a typical member of an interesting and useful class of connecting networks that has been suggested and studied by C. Clos.¹ We refer to this class as that of *three-stage Clos networks*. Such a network consists of two symmetrical outside stages of rectangular switches, with an inner stage of square switches. It is completely determined by the integer parameters m , n , r that give the switch dimensions. In one of the few outstanding contributions to the theory of connecting networks,




Fig. 1 — Three-stage Clos network $N(m, n, r)$.

Clos¹ showed that for $m \geq 2n - 1$ the network is nonblocking in the strict sense. The network defined by the parameters m, n, r will be denoted by $N(m, n, r)$.

II. SUMMARY

The following two known results about rearranging three-stage Clos connecting networks are discussed:

- i. The Slepian-Duguid theorem, which states that the network $N(m, n, r)$ is rearrangeable if and only if $m \geq n$.
- ii. The theorem of M. C. Paull, which states that if $m = n = r$, then at most $n - 1$ existing calls need be moved in $N(n, n, n)$ in order to connect an idle terminal pair.

The principal new result proven is a generalization (and possible improvement) of Paull's bound in (ii) for any m, n, r with $m \geq n$ to $r - 1$.

The Slepian-Duguid theorem is proved in Section III by an inductive method due to Duguid³ depending on the combinatorial theorem of P. Hall on distinct representatives of subsets. We discuss Paull's theorem in Section IV, but defer our simple proof of it to Section VI, which presents simple inductive proofs of various bounds on the number of calls that must be moved. All the proofs to be given depend on a "canonical reduction" procedure that consists in removing a middle switch from the network and reducing the parameters m and n by unity.

III. THE SLEPIAN-DUGUID THEOREM

The present paper is devoted to studying the property of rearrangeability for three-stage Clos networks. We shall particularly be concerned first with the possibility of rearranging calls, and later with the number of calls that must be moved. Strictly nonblocking Clos networks will not be considered except incidentally, in view of Clos's own definitive study of them.¹

Our first result is due essentially to D. Slepian,⁴ and is

Theorem 1 (Slepian-Duguid): Every three-stage Clos network with $m \geq n$ is rearrangeable.

Proof: The proof to be given is due to A. M. Duguid.^{3*} Slepian's proof was stated for the case $m = n = r$, but actually gave an explicit procedure for rearranging the existing calls so that the additional desired call could be put up. He showed for this case that at most $2n - 2$ calls must be disturbed. This bound was subsequently improved to $n - 1$ by M. C. Paull.⁵ (See Section IV.)

Duguid's proof depends on a combinatorial theorem of P. Hall, which has recently come into prominence in studies of maximal flows in networks. (See D. Gale.⁶)

Hall's Theorem: Let A be any set, and let A_1, A_2, \dots, A_r be any r subsets of A . A necessary and sufficient condition that there exist a set of distinct representatives a_1, \dots, a_r of A_1, \dots, A_r , i.e., elements a_1, \dots, a_r of A such that

$$\begin{aligned} a_i &\in A_i & i = 1, \dots, r \\ a_i &\neq a_j & \text{for } j \neq i, \end{aligned}$$

is that for each k in the range $1 \leq k \leq r$, the union of any k of the sets A_1, \dots, A_r have at least k elements.

The condition given is obviously necessary. The interest of the theorem, and our application of it, concern the sufficiency.

We proceed now to the proof of Theorem 1. It is obviously sufficient to consider only the case $m = n$. Let the inlets to the network be denoted by u_1, \dots, u_N , where $N = nr$, and let the outlets be denoted by v_1, \dots, v_N . It is sufficient to prove that every maximal assignment of inlets to outlets can be realized by a state of the network. Here "maximal" means that each inlet is to be connected to exactly one outlet, and vice versa. Such a maximal assignment is obviously equivalent to a permutation on N objects. We let $\{i \rightarrow \pi(i), i = 1, \dots, N\}$ be such a permutation; also we denote the j th inlet switch by I_j and the j th outlet switch by O_j . It is convenient to think of I_j as the set of i for which u_i is on the j th inlet switch, and of O_j as the set of i for which v_i is on the j th outlet switch.

Let K be the set of integers $\{1, 2, \dots, n\}$. We define the subsets $\{K_i, i = 1, \dots, n\}$ of K by the condition

$$K_i = \{j \mid \pi(m) \in O_j \text{ for some } m \in I_i\}.$$

* In a private communication from J. H. Déjean, the author has learned that Theorem 1 was also proved by J. LeCorre in an unpublished memorandum dated 1959.

Now let $I_{i(1)}, \dots, I_{i(k)}$ be any k of the inlet switches, and set

$$T = \bigcup_{j=1}^k K_{i(j)}.$$

Suppose that there are t distinct elements in T . Then all the kr inlets in the set

$$\bigcup_{j=1}^k I_{i(j)}$$

are assigned by $\pi(\cdot)$ to outlets from t of the outlet switches, that is, to outlets from a set of tr outlets. But two distinct inlets are not assigned to one outlet, so $t \geq k$. Thus any union of k sets among the K_i contains at least k elements.

Hence by Hall's Theorem there is a set of distinct representatives $\{k(i), i = 1, \dots, n\}$ with

$$k(i) \in K_i \quad i = 1, \dots, n$$

$$k(i) \neq k(j) \quad \text{for } i \neq j.$$

Since K contains n elements, it follows that $\{i \rightarrow k(i), i = 1, \dots, n\}$ is a permutation. However, the interpretation of the fact that $k(i) \in K_i$ is that

$$\pi(m) \in O_{k(i)} \quad \text{for some } m \in I_i.$$

In other words, to every inlet switch I_i there corresponds a unique outlet switch $O_{k(i)}$ such that $\pi(\cdot)$ maps some inlet on I_i into some outlet on $O_{k(i)}$. That is, there is a subassignment of $\pi(\cdot)$ that involves exactly one terminal on every inlet and outlet switch.

It is evident that such a subassignment can always be satisfied on a single middle switch (Fig. 1), say that numbered 1. If this subassignment is completed, that one switch is filled to capacity, and the rest of the network is essentially $N(m-1, n-1, r)$, i.e., that of Fig. 1 with the parameters m, n reduced by unity.

The theorem is clearly true for $m = n = 1$. As an hypothesis of induction assume that it is true for a given value of $m-1 (= n-1)$. The argument given above proves that it is then also true for $m (= n)$, for the induction hypothesis implies that the remainder of the assignment $\pi(\cdot)$ that was not put up on the first switch is satisfiable in the subnetwork, i.e., essentially in $N(m-1, n-1, r)$. Hence $\pi(\cdot)$ is realizable, and the theorem follows by induction on n .

IV. THE NUMBER OF CALLS THAT MUST BE MOVED: PAULL'S THEOREM

In view of the result of Slepian and Duguid that every three-stage Clos network with $m \geq n$ is rearrangeable, it is natural to ask, for a given state x of such a network, how many calls of x need actually be changed to new routes in order to put in a given call between idle terminals. Slepian's original procedure was for the case $m = n = r$, and gave the upper bound $2n - 2$ (uniformly for all states) to the number of calls that must be disturbed. That is, he showed that if $m = n = r$, then at most $2n - 2$ calls need be rearranged. By a similar but more complicated method, M. C. Paull⁵ halved this bound, proving

Theorem 2: Let $N(n, n, n)$ be a three-stage Clos network with $m = n = r$. Let x be an arbitrary state of this network. The largest number of calls in progress in x that must be rerouted in order to connect an idle pair of terminals is $n - 1$; there exist states which achieve this bound.

Since Paull's proof was involved, we have looked for and found simpler ways of proving and extending his result. In Section VI we give a simple inductive proof; the argument to be given, of course, also provides a proof of the Slepian-Duguid theorem not depending on the Hall combinatorial result used in Section III.

V. SOME FORMAL PRELIMINARIES

In order to state and prove the rest of our results, it is useful, and indeed necessary, to introduce a systematic notation. Such a notation has been described and used in a previous paper² by the author; the notation to be used is a consistent extension of this.

The set of inlets of a network is denoted by I , and that of outlets by Ω . The set of possible states of a connecting network is denoted by S . For a three-stage Clos network, S consists of all the ways of connecting a set of inlets to as many outlets by disjoint chains (paths) through an inlet switch, a middle switch, and an outlet switch. (See Fig. 1.) States of the network may then be thought of as sets of such chains. Variables x, y, z, \dots , at the end of the alphabet, range over states from S .

A terminal pair $(u, v) \in I \times \Omega$ (with u an inlet and v an outlet) is called *idle in state x* if neither u nor v is an endpoint of a chain belonging to x . A *call c* is a unit subset $c = \{(u, v)\} \subset I \times \Omega$; c is *new in a state x* if (u, v) is idle in x . The *assignment $\gamma(x)$* realized by x is the union of all calls $c = \{(u, v)\}$ such that x contains a chain from u to v . If a is an assignment, $\gamma^{-1}(a)$ is the set of all states realizing a . The cardinality of a set X is denoted by $|X|$. The states $x \in S$ are partially ordered by inclusion \leq in a natural way.

A distance between states can be defined as

$$\begin{aligned}\delta(x, y) &= |x \Delta y|, \\ &= \text{the number of calls that would have to be added, removed,} \\ &\quad \text{or rerouted to change } x \text{ into } y,\end{aligned}$$

where Δ is symmetric difference. The distance of a state x from a set X of states is defined in the usual way as

$$\delta(x, X) = \min_{y \in X} \delta(x, y).$$

A call c new in a state x is blocked in x if there is no state $y > x$ such that $\gamma(y) = \gamma(x) \cup c$. A state x is nonblocking if no call new in x is blocked in x . The set of nonblocking states is denoted by B' . For any call c , the set of states x in which c is both new and not blocked is designated B_c' .

For a three-stage Clos network $N(m, n, r)$ with $m \geq n$ we define

$$\begin{aligned}\varphi_x(m, n, r) &= \max_{c \text{ new in } x} \delta(x, \gamma^{-1}(\gamma(x) \cup c)) - 1 \\ &= \max_{c \text{ new in } x} \delta(x, \gamma^{-1}(\gamma(x)) \cap B_c') \\ &= \max_{c \text{ new in } x} \min_{y \in \gamma^{-1}(\gamma(x)) \cap B_c'} \delta(x, y) \\ &= \text{the maximum number of calls that must be re-} \\ &\quad \text{routed in order to put up a call } c \text{ new in } x.\end{aligned}$$

We also set

$$\varphi(m, n, r) = \max_{x \in S} \varphi_x(m, n, r).$$

In this last definition, it is assumed that S is the set of states determined by the parameters m, n, r in Fig. 1.

In the notation introduced above, the Slepian-Duguid Theorem guarantees that for $m \geq n$ and c new in x

$$\begin{aligned}\gamma^{-1}(\gamma(x) \cup c) &\neq 0, \\ \gamma^{-1}(\gamma(x)) \cap B_c' &\neq 0,\end{aligned}$$

and Paull's Theorem may be cast as stating that

$$\varphi(n, n, n) = n - 1.$$

VI. THE NUMBER OF CALLS THAT MUST BE MOVED: NEW RESULTS

We now present some new methods for studying the number of calls that must be moved; these yield extensions of results of D. Slepian⁴ and M. C. Paull.⁵

Theorem 3: $\varphi(2, 2, r) \leq 2r - 2$.

Proof: Suppose that a blocked new call between input switch I_1 , and output switch O_1 is to be put in when the network is in a state x . Consider any sequence c_1, \dots, c_k of existing calls of x with the properties

i. Either c_1 is on I_1 , c_1 and c_2 are the same outlet switch, \dots ,

c_i and c_{i+1} are on the same outlet switch, i odd, $i < k$

c_i and c_{i+1} are on the same inlet switch, i even, $i < k$,

or c_1 is on O_1 , c_1 and c_2 are on the same inlet switch, \dots ,

c_i and c_{i+1} are on the same inlet switch, i odd, $i < k$

c_i and c_{i+1} are on the same outlet switch, i even, $i < k$.

ii. c_k is the only call on some outer switch. Since neither I_1 nor O_1 is full, the largest k for which such a sequence exists is $2r - 2$. The reader can verify that a possible strategy for rearranging existing calls of x so as to put in an I_1 - O_1 call is to take each call of the sequence c_1, \dots, c_k and reverse its route, i.e., make it go through the *other* middle switch than the one it presently uses. Thus for all x

$$\varphi_x(2, 2, r) \leq 2r - 2.$$

Let x be a state of $N(m, n, r)$, and let M be a particular middle switch. A *canonical reduction* of x with respect to M will consist of

i. removing M ,

ii. on each outer switch that has a call routed via M , removing the link, crosspoints, and terminals associated with that call,

iii. on each outer switch that has an idle link to M , removing the link, the crosspoints associated therewith, and one arbitrarily chosen idle terminal.

It is easily seen that a canonical reduction of a state x of $N(m, n, r)$ leads to a state of $N(m - 1, n - 1, r)$.

Theorem 4: $\varphi(n, n, r) \leq 2r - 2$.

Proof: By Theorem 3, the result holds for $n = 2$, so assume it for a given value of $n - 1 \geq 2$, and try to rearrange a given state x of $N(n, n, r)$ so as to put in a new blocked call from I_1 to O_1 .

Case 1: There is a middle switch M with both an I_1 and an O_1 call on it. Perform a canonical reduction of the state x with respect to M . This yields a state of $N(n - 1, n - 1, r)$, for which the result holds.

Case 2: No middle switch has both an I_1 and an O_1 call on it. Since the call to be put in is blocked, it must be true that

$$\#(\text{idle links out of } I_1) + \#(\text{idle links out of } O_1) = n$$

and hence

$$\max\{\# \text{ (idle out of } I_1), \# \text{ (idle out of } O_1)\} > 1.$$

Suppose that $\# \text{ (idle out of } I_1) > 1$. There is a middle switch M with an idle link to I_1 , and a busy link to O_1 . Perform a canonical reduction of x with respect to M , yielding a state of $N(n-1, n-1, r)$ in which each of I_1, O_1 still has an idle terminal.

A refinement of this method suggested by M. C. Paull will halve the last two bounds. We prove

Theorem 5: $\varphi(2, 2, r) \leq r - 1$. ($r \geq 2$)

Proof: The result is true for $r = 2$, since in that case the network has only one blocking state (see Fig. 2), and both blocked calls can be unblocked by changing the route of one ($= r - 1$) existing call.

Let us assume as an hypothesis of induction that the theorem holds for some value of $r - 1 \geq 2$, and in $N(2, 2, r)$ attempt to put up a blocked new call c between input switch I_1 and output switch O_1 . Since c is new and blocked, there must be an idle and a busy link on both of I_1 and O_1 , and each of the busy links must pass through a different middle switch. Let c_1 be the call on I_1 , and c_2 be the call on O_1 . We may suppose without loss of generality that c_1 is a call from I_1 to O_2 , while c_2 is a call from I_2 to O_1 .

Case 1: I_2 has only one call on it, viz., c_2 . Move c_2 to the other middle switch (see Fig. 3).

Case 2: I_2 has two calls on it. Remove both c_1 and c_2 , so that I_1 and O_1 become empty. Consider now the state x of the subnetwork of parameter $r - 1$ obtained by removing I_1 and O_1 and reducing the dimension of the two square middle switches by unity to $r - 1$. Each of I_2 and O_2 has at least one idle terminal in x , since c_1 and c_2 were removed. Hence by the hypothesis of induction the subnetwork can be rearranged so as to put in a call from I_2 to O_2 while disturbing at most $r - 2$ existing calls. If the I_2 - O_2 path thus provided is via M_1 then c_1 and c_2 can be replaced as in




Fig. 2 — Network with only one blocking state ($r = 2$).


Fig. 3 — I_2 with one call, c_2 .

Fig. 4. This leaves a path for the new I_1 - O_1 call c via M_2 , and shows that it was never necessary to move c_2 , and that hence at most $r - 1$ calls were disturbed. If the I_2 - O_2 path provided by rearranging the subnetwork is via M_2 , then c_1 and c_2 can be replaced as in Fig. 5. This leaves a path for c via M_1 , and shows that c_1 did not really have to be moved, so that at most $r - 1$ calls were disturbed.


Theorem 6: $\varphi(n, n, r) \leq r - 1$.


Proof: The result is true for $n = 2$. Assume that the theorem is true for a given value of $n - 1 \geq 2$, and seek to rearrange a state x of $N(n, n, r)$ so as to put in a new call blocked in x between I_1 and O_1 . The theorem follows by induction on n by distinguishing two cases as in Theorem 4, and using a canonical reduction of x .

Theorem 7: For $m - 1 \geq n$,

$$\varphi(m, n, r) \leq \varphi(m - 1, n, r).$$

Proof: This is almost obvious. Remove any middle switch M of $N(m, n, r)$ and make all terminals on which there were calls routed via M idle. This gives a state of $N(m - 1, n, r)$; in this state the desired call

Fig. 4 — Calls c_1 and c_2 over path via M_1 .

Fig. 5 — Calls c_1 and c_2 over path via M_2 .

can be put in by rearranging at most $\varphi(m-1, n, r)$ existing calls. Now replace M and the calls that were routed through it.

M. C. Paull³ has conjectured that if $r = n$, then

$$\varphi(m, n, n) \leq 2n - 1 - m.$$

This bound agrees with Theorem 2 if $m = n$, and with Clos' results on nonblocking networks if $m = 2n - 1$. Paull has proved the result for $m = 2n - 2$. However, no proof of the full conjecture has been found. It is tempting to try the stronger conjecture that

$$\varphi(m, n, r) \leq 2n - 1 - m$$

for any m, n , and r . This can be disproved by the counterexample shown in Fig. 6. There is no way of connecting I_1 to O_5 without moving a call




Fig. 6 — Network showing that I_1 and O_5 cannot be connected without moving a call on one of I_1, O_5 .

on one of I_1, O_6 . However, all possible alternative routes for these calls are pre-empted, so at least two calls must be moved.

VII. ACKNOWLEDGMENT

The author is indebted to M. C. Paull for helpful comments, and for suggestions that led to Theorems 5 and 6.

REFERENCES

1. Clos, C., A Study of Non-Blocking Switching Networks, B.S.T.J., **32** (1953), pp. 406-424.
2. Beneš, V. E., Algebraic and Topological Properties of Connecting Networks, B.S.T.J., **41** (1962), pp. 1249-1274.
3. Duguid, A. M., Structural Properties of Switching Networks, Brown University, Progress Report BTL-7, 1959.
4. Slepian, D., Two Theorems on a Particular Crossbar Switching Network, unpublished manuscript, 1952.
5. Paull, M. C., Reswitching of Connection Networks, B.S.T.J., **41** (1962), pp. 833-855.
6. Gale, D., A Theorem on Flows in Networks, Pacific Journal of Mathematics, **7** (1957), pp. 1073-1082.